

A Note on the Simple Exponential Smooth Non-optimal Predictor, the Order-up-to Policy and How to Set a Proper Bullwhip Effect

Erland Nielsen

► **To cite this version:**

Erland Nielsen. A Note on the Simple Exponential Smooth Non-optimal Predictor, the Order-up-to Policy and How to Set a Proper Bullwhip Effect. 19th Advances in Production Management Systems (APMS), Sep 2012, Rhodes, Greece. pp.622-629, 10.1007/978-3-642-40361-3_79 . hal-01470676

HAL Id: hal-01470676

<https://hal.inria.fr/hal-01470676>

Submitted on 17 Feb 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



A note on the simple exponential smooth non-optimal predictor, the order-up-to policy and how to set a proper bullwhip effect

Erland Hejn Nielsen*

CORAL - Centre for Operations Research Applications in Logistics,
Aarhus University, Business and Social Sciences,
Department of Economics and Business,
Fuglesangs Allé 4, DK-8210 Aarhus V, Denmark
Phone: (+45) 8716-4903
{ehn}@asb.dk
<http://au.dk/en/ehn@asb.dk>

Abstract. *The literature concerning the bullwhip effect is mostly focused on determining expressions for the theoretical bullwhip measure given specific theoretical system setups, whereas it must also be of interest to deal with the problem of how in fact to make a proper choice as to a sensible bullwhip level. Such a management approach to the bullwhip phenomenon has to be of quite some importance, as the bullwhip effect on the one side definitely is a system malfunction, but on the other also an effect the size of which common intuition tells us should be possible to control. The control is based on a decision as to what variation in demand should be locally absorbed and what variation should be passed on upstream. This paper will focus on design aspects of a bullwhip control policy in order to decide on sensible trade-offs between the bullwhip level and the local inventory variability.*

Keywords: The bullwhip effect, inventory variance, non-stationary demand, unobserved parameters

* Acknowledgment: Initial draft versions of this work have been supported by grant no. 275-07-0094 from the Danish Social Science Research Council.

1 Introduction

When digging down into the vast body of literature dealing with the bullwhip effect, it seems quite possible to get a fairly good understanding of the phenomenon on the generic level, but when it comes to the actual managerial choice of an appropriate bullwhip value, the story seems to be less decisive leaving us with an undetermined trade-off between the variation in demand that is absorbed locally into the inventory and the variation in demand that is propagated upstream in the supply chain. Somehow, such trade-off has always been at the core topic when dealing with standard production and inventory control [7] under the headings of production smoothing or production leveling. The role of forecasting in much classical production and inventory theory is often a combined one in that the forecasting mechanism also handles the production smoothing upstream. The simple exponential smoothing mechanism is one of the most popular schemes in that respect and the smoothing coefficient is by experience recommended to be held within the interval 0.1 to 0.3 typically.

Any supply chain collaboration or coordination aims at balancing production smoothing with inventory control at each tier. A basic intuition tells us that the less smoothing locally the more variation is sent upstream the supply chain. So which are the controlling handles when talking about the bullwhip effect, and how should these handles be set? And do we really have this inverse relation between variation absorbed locally and variation exported upstream?

Let us consider as a starting point the discussion given by Stephen C. Graves [5], where he develops the bullwhip formula $BW = (1 + (1 - \theta)L)^2$ derived conditionally on the optimal predictor given that the demand DGP (Data Generating Process) is IMA(1,1). The optimal predictor in this case is the simple exponential smooth mechanism with the smoothing coefficient $\alpha = 1 - \theta$. In the Graves setup we are then left with the parameters θ and L , both of which are not really controllable in the short run, but given by the structure of the supply chain.

If we want to get access to the non-optimal exponential smoothing coefficient values, we need to consider bullwhip measures unconditionally. However, the unconditional variance of a non-stationary variate, as for instance a random walk (IMA(1,1) where $\theta = 0$), is infinite.

We therefore have to redefine the bullwhip measure in this case to be focused only on the amplifying signal effect on top of the (stochastic) trend movements. The traditional way to eliminate the (stochastic) trend in the classical ARIMA [1] approach is to consider the differenced data instead.

Henceforth, the bullwhip measure that will be used is then defined as

$$BW = \frac{Var(\Delta q)}{Var(\Delta D)} \quad (1)$$

where q denotes the upstream ordering and D is the demand.

The paper will be organized in the following way: In section 2 the simulation model used in this work is described and commented upon. Section 3 is devoted to mapping out the bullwhip measure as well as the inventory variation effect as

a function of the simple exponential smoothing factor. Section 4 is dealing with the trade-off between absorbing demand variation in the local tier inventory and passing the variation upstream as bullwhip. Finally, section 5 presents a few concluding remarks.

2 The supply chain simulation model

The model used in order to simulate the actual bullwhip values follows tradition. Customer demand as seen by the retailer is assumed to be well described by a non-stationary time series model of the ARIMA(0,1,1) type process

$$\Delta D_t = \varepsilon_t - \theta \cdot \varepsilon_{t-1} \quad \text{where } \varepsilon_t \sim \text{i.i.d. } N(0, \sigma_\varepsilon^2) \quad (2)$$

This process is covariance stationary for any value of θ , however, the MA(1)-part of the model is only invertible for values of the parameter θ between -1 and 1. Whenever $\theta=0$, the IMA(1,1) alias ARIMA(0,1,1) is equivalent to the dynamic process which is normally denoted "a random walk". As we are facing a delivery lead time L typically of a non-negligible magnitude, it is necessary to have a sound idea of the cumulative amount of demand that potentially could be executed during such a period. A measure y_t of such a lead time demand could be formalized as follows

$$\begin{aligned} \hat{D}_t(1) &= \hat{D}_{t+1} = S_t = \alpha \cdot D_t + (1 - \alpha) \cdot S_{t-1} \\ \hat{D}_t(L) &= L \cdot \hat{D}_t(1) \\ \hat{\sigma}_{et}^2(L) &= \hat{\sigma}_\varepsilon^2 \cdot \{1 + (L - 1) \cdot \alpha^2\} \\ y_t &= \hat{D}_t(L) + z \cdot \hat{\sigma}_{et}^L \end{aligned} \quad (3)$$

where $-1 < \theta < 1$ and $0 \leq \alpha \leq 1$. This setup represents the optimal predictor setup if $\alpha = 1 - \theta$, in terms of minimum mean square error forecast, for the IMA(1,1) process and the z -value (≥ 0) controls to some extent the "worst case" dependence as a function of the variance of the stochastic innovations of the IMA(1,1) demand process induced on the cumulative future demand.

The determination of the size and timing of orders q_t that are placed by the retailer upstream at the manufacturer can be controlled by a variety of ordering mechanisms [4], however, if we assume an order-up-to "base-stock"-policy as is often encountered in the literature (see for instance [2][3][5]) and which are basically of a make-to-order nature, q_t can be expressed as follows

$$q_t = y_t - y_{t-1} + D_t \quad (4)$$

Inventory is then simply a matter of elementary "stock-flow" bookkeeping

$$I_t = I_{t-1} - D_t + q_{t-L} \quad (5)$$

The simulation setup is a periodic time setup and 50,000 time steps (approx. steady-state) are being simulated for each combination of $\alpha = 0$ to 1 in steps by 0.02 and $\theta = 0$ to 0.9 in steps by 0.1. The optimal predictor situations $\alpha = 1 - \theta$ are singled out specifically. Demand is simulated with $D(init) = 0$ and error term Std.dev=10 and afterwards level shifted in order to be exactly non-negative. Inventory is also controlled for non-negativity. $I(init) = 5,000$ seems to do the job. z is set equal to zero throughout this paper. The actual computations are programmed and performed in the Gnu-R statistical system[6].

3 Mapping out the bullwhip and inventory variance effect

The variance components that will be dealt with in this section are $Var(\Delta D)$, $Var(\Delta q)$ and $Var(I)$. In order for these computations to be meaningful ΔD , Δq and I all have to be covariance stationary. ΔD complies by construction, whereas Δq and I are checked manually. For $\alpha = 0$ the inventory dynamics are defined by $I_t = I_{t-1} - D_t + D_{t-L}$ and here it seems that non-stationarity takes over resulting in $Var(I)$ approaching infinity. For $\alpha = 0$ the bullwhip effect is trivially neutral, that is $BW = 1$. However, the full BW -map turns out as follows.




Fig. 1. Bullwhip measures as a function of α (Alpha) given $\theta = 0$ to 1 in steps of size 0.1 (the curves); Demand as IMA(1,1) with error term Std.Dev=10 and a lead time of $L=3$. Each bullwhip curve is θ -specific and the curve dots are indicating the $\alpha = 1 - \theta$ situations, thereby implicitly identifying the respective curves' specific θ value.

The message from the obtained BW -map is fairly simple. Irrespective the actual state of the DGP for the demand and a fixed lead time L , a smaller value of α results in a smaller bullwhip effect. Furthermore, deviations from the optimal predictor setting are almost of no consequence for small α values ($\alpha \leq 0.5$) and of some consequence otherwise. This is definitely interesting from a production leveling perspective. It can furthermore be noted that the curvature of the BW -curves based on the first time differenced data is the same as the Graves ([5]) BW measure. The Graves BW measure is, however, consistently lower having a value of 16 instead of 25 for $\alpha = 1$ and $\theta = 0$. This is not surprising in that the stochasticity that works through the simple exponential smoothing mechanism has been taken completely out in the Graves setup due to the conditioning.

So being able to control the bullwhip effect by the simple exponential smoothing factor, the mirror effect on the local inventory dynamics has to be observed. The full $Var(I)$ -map turns out as follows:




Fig. 2. Inventory variance (VAR-INVENTORY(i,)) as a function of α (Alpha) given $\theta = 0$ to 1 in steps of size 0.1 (the curves); Demand as IMA(1,1) with error term Std.Dev=10 and a lead time of $L=3$. Each variance curve is θ -specific and the curve dots are indicating the $\alpha = 1 - \theta$ situations, thereby implicitly identifying the respective curves' specific θ value.

The message from the obtained $Var(I)$ -map is a little more complicated than for the BW -map in that choosing a too low α easily results in an extremely large inventory variance implying that the local inventory is virtually out of control. Take as an example the top curve on the $Var(I)$ -map with the dot

corresponding to $\alpha = 1$ and thereby identifying the $\theta = 0$ curve. In this case the demand is a pure random walk and the optimal one-step-ahead predictor is the so-called naïve forecast. Even if demand is characterized by a stochastic trend, it is quite impossible to even predict whether demand goes up or down in the next period. The inventory variance is contrary to the BW measure dependent on both scale and level of noise of the demand process. In the analyzed case the optimal predictor situation implies an inventory variance of approximately 1,300, that is a standard deviation of ± 36 units. If in this case we decided to choose a non-optimal but lower value of α for example 0.5, the resulting inventory variance would be approximately 1,700, that is a standard deviation of around ± 41 units. Maybe not so drastic an increase, but if α is decreased to 0.3, the variance increases to 2,200, which then is a standard deviation of around ± 47 units. And finally if α is decreased to 0.1, the variance increases to 5,200, which then is a standard deviation of around ± 72 units. Compared to the standard deviation on demand ± 10 units it is quite impressive though.

The above given example constitutes the worst case situation, but if we now turn our reasoning around and assume that we are not quite sure of the actual θ value and tend to go for a production smoothing value of α say 0.2, then we may expect the inventory variance to be somewhere between 500 and 3000, translated into standard deviations between 22 and 55. Moving up to $\alpha = 0.5$, the standard deviations must be expected to lie between 24 and 45, which is quite a reduction in the uncertainty range. On the other hand going from $\alpha = 0.2$ to $\alpha = 0.5$ the bullwhip changes from being within the interval [2.6; 2.8] to being within the interval [7.0; 8.5].

If we know the actual θ value from the demand DGP, we can observe that an optimal choice of $\alpha = 1 - \theta$ simply results in the minimal obtainable local inventory variance. This is quite a nice property, but it implies that the optimal prediction setup in the present case simply chooses a BW effect that is based on a maximum transference of upstream variability. This is maybe not so nice a property if the BW effect really hurts upstream. Anyway, it will clearly depend on specific trade-off considerations in given situations.

4 How to set a proper bullwhip effect

In order to decide about the choice of a "proper" bullwhip effect, here given the system's "malfunctioning" with respect to ordering, there are obviously two very different situations. One where the demand DGP parameters have been determined successfully and one where the actual parameter settings are unknown, but only known to be an IMA(1,1) structure.

The figure below is constructed based on the observation that α and BW are almost one-to-one at least for α values below 0.5. For α values above 0.5 a midpoint BW value is used.

The curves "UB Std.Dev(I)" and "LB Std.Dev(I)" represent the decision information where θ is unknown, whereas if θ is known, the "OPT Std.Dev(I)" curve represents the decision information.

Clearly, in the case of a known θ there is not much of a trade-off, but approximately a one-to-one strictly increasing relationship between BW and $Var(I)$. But then again, there is not much of a managerial choice present in this situation.




Fig. 3. BW against Std.Dev(I)-ranges.

More interesting is the situation where only the structural form is known, but none of its parameters. It must, however, be remembered that besides not knowing θ we then also do not know the size of the error of the innovation term of the demand DGP, and so any reasoning based on figure 3 is clearly only indicative as it is simulated based on a specific error term Std.Dev=10.

Nevertheless, we now have a trade-off situation with respect to the "UB Std.Dev(I)". In a worst case sense the trade-off along this curve represents a lower expected inventory variance traded against a higher BW effect and vice versa.

But maybe an interpretation that connects the range of uncertainty, to a choice of a BW effect is the really interesting angle on the subject of setting a proper BW effect. Small BW values correspond well to small α settings, but leave the local inventory with a huge uncertainty as to the actual dynamic variability. This uncertainty effect is certainly narrowed down for increasing BW values.

The question is now really - how do we value risk with respect to local inventory variance against upstream exported and amplified demand variation as expressed by the dynamics of the ordering? Clearly, this is closely related to the worst case ("UB Std.Dev(I)") trade-off against BW , but the risk element focuses more clearly on the element of gambling that is virtually always present.

5 Concluding remarks

It is quite obvious that this work has not really given a complete answer to the posed question of how to set a proper BW effect. Still a few moments to remember have come by in that just following classical production and inventory control wisdom and setting the smoothing factor low, somewhere between 0.1 and 0.3, might be a risky business for the local inventory even if it produces a small BW effect upstream. If it was the bullwhip effect alone, α should trivially be set to zero. Now if θ is known and is low, which means that the demand process is closer to a pure random walk, then a non-optimal low α choice is really bad. We are then simply close to the "UB Std.Dev(I)" curve values. Pressing α towards zero in this situation makes things even worse in that the inventory variance goes towards infinity. So, at last and no surprise, the concrete choice of upstream bullwhip effect does in the end depend on precisely how costly unforeseen inventory and/or production activity variation is at the individual tiers in the supply chain. But the intuition that somehow there must always be a negative trade-off between the variance absorbed locally and the variance transferred up-stream seems also to be supported in the specific case studied, where demand is non-stationary and the order mechanism is of the order-up-to type, at least viewed from a certain perspective.

References

1. Box G., Jenkins G.: Time series analysis: Forecasting and control. Holden-Day, San Francisco (1970)
2. Chen, F., Drezner, Z., Ryan, J. K., Simchi-Levy, D.: Quantifying the Bullwhip Effect in a Simple Supply Chain: The Impact of Forecasting, Lead Times, and Information. *J. Management Science*. 46,(3) 346-443 (2000)
3. Chen, F., Ryan, J. K., Simchi-Levy, D.: The impact of exponential smoothing forecasts on the Bullwhip Effect. *J. Naval Research Logistics*. 47, 269-286 (2000)
4. Disney, S. M., Lambrecht, M. R.: On Replenishment Rules, Forecasting and the Bullwhip Effect in Supply Chains. *Foundations and Trends in Technology, Information and Operations Management*. 2,(1), 1-80 (2007)
5. Graves, S.C.: A Single-Item Inventory Model for a Nonstationary Demand Process. *J. Manufacturing & Service Operation Management*. 1,(1) 50-61 (1999)
6. The R Project for Statistical Computing, <http://www.r-project.org/>
7. Silver E. A., Pyke D. F. and Peterson R.: Inventory Management and Production Planning and Scheduling. John Wiley & Sons, Inc. New York (1998)